

DOCUMENTED
BRIEFING

A Simple Game-Theoretic
Approach to Suppression
of Enemy Defenses and
Other Time Critical Target
Analyses

Thomas Hamilton, Richard Mesic



PROJECT AIR FORCE

20041108 072

BEST AVAILABLE COPY

A Simple Game-Theoretic Approach to Suppression of Enemy Defenses and Other Time Critical Target Analyses

Thomas Hamilton, Richard Mesic

DB-385-AF

August 2004

Prepared for the United States Air Force

Approved for public release; distribution unlimited



The research reported here was sponsored by the United States Air Force under Contract F49642-01-C-0003. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

ISBN: 0-8330-3259-3

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2004 RAND Corporation

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2004 by the RAND Corporation
1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516
RAND URL: <http://www.rand.org/>
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

ISBN: 0833032593

Technical rpt #: RAND/DB-385-AF

Order No./Price: \$15.00 paperback

Personal name: Hamilton, Thomas.

Title: A simple game-theoretic approach to suppression of enemy defenses and other time critical target analyses / Thomas Hamilton, Richard Mesic.

Electronic access: Online access <http://www.rand.org/publications/DB/DB385/>

Publication info: Santa Monica, CA : RAND, 2004.

Physical description: ix, 53 p. : ill. (some col.) ; 28 cm.

Note: "August 2004."

Note: "Project Air Force."

Note: Includes bibliographical references.

Abstract: The effectiveness of attacks on time critical targets (suppression of enemy air defenses, interdiction, and anti-theater ballistic missile missions) often depends on decisions made by the adversary. Game theory is a way to study likely changes in enemy behavior resulting from various attack capabilities and goals. Engagement-level combat is treated as a two-player game in which each player is free to choose its strategy. The response an intelligent opponent is likely to make to differing levels of threat capability is critical to understanding and measuring the capability necessary to induce the enemy to follow a preferred course of action. Enemy willingness to engage is an important factor. If the enemy decides not to launch missiles or move ground vehicles, it has become paralyzed, in itself a worthy goal. The emphasis in the study is on the choice of strategies in realistic military situations; all can be analyzed with straightforward mathematics. Finally, the authors discuss situations in which the two sides have different views of the duration of the conflict or the appropriate measures of effectiveness. It is a great advantage to a combatant to know the opponent's real objectives.

Other forms: Also available on the internet via WWW in PDF format.

Ctrct/Grnt/Proj/Task: Air Force. F49642-01-C-0003 . PA002 3953. PA002 4961

Subject: War games--Planning--Simulation methods.

Subject: War games--Evaluation--Simulation methods.

Subject: Special operations (Military science)--United States--Simulation methods.

DTIC descriptor: War games.

DTIC descriptor: Air Force.

DTIC descriptor: Strategy.

DTIC descriptor: Warfare.

Personal name: Mesic, Richard, 1943-

Corporate name: Project Air Force (U.S.). Manpower, Personnel, and Training Program.

Corporate name: Rand Corporation.

Corporate name: United States. Air Force.

Series: Documented briefing (Rand Corporation) ; 385

Research unit: Project Air Force

Distribution code: 3

Subject bibliography: MILITARY STRATEGY AND TACTICS

Subject bibliography: GAMING

Distribution date: 20040909

PREFACE

This documented briefing uses simple game theory to relate U.S. Air Force operations to the likely effects of such operations on enemy behavior. This approach provides an analytically sound basis for the quantitative study of the military operations, where the term "effects" has a broad, but rigorous, meaning.

We analyze attacks on time critical targets, with emphasis on suppression of enemy air defenses (SEAD), interdiction, and anti-theater ballistic missile missions. In all these cases, enemy willingness to engage is an important factor. If the enemy decides not to launch missiles or move ground vehicles, no attack is possible, but paralyzing the enemy by forcing such decisions is in itself a worthy goal.

We analyze engagement-level combat as a two-player game in which each player is free to choose from a variety of strategies. Game theory provides quantitative tools to calculate and evaluate optimal and suboptimal strategies for each player. Critical to this problem is understanding the response an intelligent opponent is likely to make to differing levels of threat capability. In this way we can measure what capability is necessary to induce our enemy to follow a preferred course of action.

Throughout, our emphasis is on the choice of strategies in realistic military situations. A wide variety of situations can be usefully analyzed with straightforward mathematics.

Finally, we discuss situations in which the two sides have different views of the duration of the conflict or the appropriate measures of effectiveness. This discussion emphasizes the great advantage to a combatant of knowing the opponent's real objectives.

This research took place in the Aerospace Force Development Program of RAND Project AIR FORCE. The documented briefing expands the game theoretic approach used in *Dynamic Command and Control: Focus on Time Critical Targets*, MR-1437-AF (government publication; not releasable to the general public). The work was sponsored by the Air Force director of Intelligence, Surveillance, and Reconnaissance (AF/XOI), and the Air Force director for Command and Control (AF/XOC).

This research should be of interest to Air Force and DoD personnel involved in developing and fielding enhanced command and control (C2) and battle management (BM) systems, and to operators developing improved C2 and BM tactics, techniques, and procedures for force-level and unit-level operations.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the U.S. Air Force's federally funded research and development center for studies and analyses. PAF provides the Air Force with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future aerospace forces. Research is performed in four programs: Aerospace Force Development; Manpower, Personnel, and Training; Resource Management; and Strategy and Doctrine.

Additional information about PAF is available on our web site at <http://www.rand.org/paf>.

SUMMARY

When planning operations against time-critical targets (TCTs), commanders typically think about how much capability they need to kill enemy forces. However, they also consider how their strategies will affect the enemy's behavior. TCT operations include suppression of enemy air defenses (SEAD), interdiction of moving forces, and attacks against theater ballistic missiles (TBMs). Convincing an enemy not to fire surface-to-air missiles (SAMs), not to move his forces, or not to launch TBMs is an equally good outcome as actually defeating his capabilities. This approach is what military analysts today refer to as "effects-based operations."

Traditional methods of planning TCT operations are not adequate to analyze these effects. Planners typically seek to determine how much capability the United States would need in order to defeat an enemy who exerts maximum effort. This approach was valid during the Cold War when an enemy faced with a nuclear air strike was expected to activate all of his air defenses at the outset of the conflict. However, in the types of military campaigns that the United States faces today, an intelligent enemy is not expected to pursue a fixed, highly aggressive strategy. Rather, he is expected to adjust his behavior in response to U.S. capabilities and actions. For example, an opponent faced with air strikes will consider the value of his SAMs, the value of the target under attack, and the level of resources U.S. forces are devoting to SEAD before he decides how much air resistance to mount. This behavior has been seen in recent conflicts such as Operation Allied Force (Kosovo) and Operation Desert Storm (Iraq). Under certain conditions, an intelligent enemy might conclude that his best option is not to fire SAMs at all. If military planners can understand how these factors influence an enemy's decisionmaking, then they can determine what capabilities and strategies the United States would need in order to paralyze an opponent.

Researchers within RAND Project AIR FORCE have developed a method that military planners can use to analyze the effects of U.S. strategy and capabilities on the enemy in TCT operations. The method uses the mathematical techniques of game theory to understand how each side of a military conflict influences the other's decisionmaking and how one side can compel the other to follow a preferred course of action.

GAME THEORY PROVIDES A METHOD FOR ANALYZING TCT OPERATIONS

Game theory was developed in the 1920s as a way of using mathematics to model human decisionmaking in competitive situations. It is ideally suited for analyzing military situations. A military conflict can be represented as a two-player game in which the interests of both sides are diametrically opposed. If one side wins, then the other side must lose. Both sides are free to choose between a variety of moves and to adjust

their strategy over the course of several turns. The advantage of game theory is that analysts do not know in advance what the enemy will do. Thus, they are forced to account for the realistic situation in which an intelligent enemy may decide for himself what his best move is. Analysts draw conclusions about the factors that compel each side to adopt a given strategy.

Military planners can apply these principles to TCT operations through game theoretic analysis. The method consists of the following steps:

- **Determine the tactical options available to each side.** Analysts assume that each opponent can choose between a range of possible actions. For example, in a simple SEAD operation, the attacker (called "Blue") can choose to fly a strike aircraft, such as an F-16 CG, or to fly a SEAD aircraft, such as an F-16 CJ. The strike aircraft attempts to get past enemy air defenses and to strike a target on the ground. The SEAD aircraft detects radio emissions from active SAM radars and attempts to destroy them. The defender (called "Red") can choose to activate his SAM radars or to leave them inactive. Analysts may construct games in which each side has more than two options. The arithmetic is more complex, but the principles are the same. (See pages 3-4.)
- **Assign a numerical value to each possible outcome.** Commanders in the field routinely make value judgments about the strength of their capabilities, the probability that their weapons will hit a target, and the potential gain or loss of a particular engagement. Analysts can represent these judgments in mathematical terms by determining the measures of effectiveness (MoEs) for each capability. For example, Figure S.1 shows the MoEs in the simple SEAD game from Blue's point of view. Analysts judge that a Blue strike aircraft is worth five points. Thus, if Blue flies a strike aircraft and Red activates his SAMs, then the aircraft is likely to be shot down and Blue will lose five points. If Blue flies a strike aircraft and

	Red SAM engages	Red SAM does not engage
Blue flies strike	-5	1
Blue flies SEAD	5	0

Figure S.1—Analysts Assign Notional MoEs to Each Outcome

Red does nothing, then Blue will win one point for moving past enemy defenses. Conversely, analysts judge that a Red SAM is also worth five points. If Blue flies a SEAD aircraft and Red SAMs engage, then Blue will win five points for likely hitting a SAM. If Red does not activate its SAMs, then Blue will gain nothing from the encounter. (See page 4.)

- **Calculate all possible strategies and their outcomes.** Intelligent opponents vary their tactics in order to appear unpredictable to the enemy. Thus a combatant's overall strategy is determined by how often he chooses one tactical option over another. Analysts may calculate all possible strategies and the mathematical outcomes of each strategy. This information may be presented using a graph such as Figure S.2. The MoEs from Figure S.1 are reproduced to the right of the graph. The x-axis represents the percentage of the time that Red chooses not to activate his SAMs. The y-axis represents the percentage of the time that Blue chooses to fly a strike aircraft. The shaded bands represent the outcome of each combination from Blue's point of view. Using this chart, analysts can understand how each combatant may alter his strategy based on the other side's actions. For example, if Blue perceives that Red withholds his SAMs most of the time, then Blue may choose to fly more strike aircraft than SEAD (upper right corner). This is a good outcome for Blue because his strike aircraft can continue on to strike enemy targets. However, Red may turn the engagement in his favor by activating his SAMs more often (upper left corner). This is a bad outcome for Blue because he will lose a high percentage of strike aircraft. Blue may reverse the outcome again by increasing the proportion of SEAD aircraft (lower left corner). Based on this kind




Figure S.2—Analysts Calculate the Results of Each Combination of Strategies

of analysis, planners can anticipate how an enemy is likely to respond to a given strategy. (See pages 5-6.)

- **Find each side's optimum strategy.** The optimum strategy is the best strategy that each side can pursue without the other side being able to turn the engagement in his favor. Experience teaches that as opponents in a game adjust to each other's actions, each player will eventually settle on his optimum strategy, also called the "equilibrium solution." In military terms, the optimum strategy is not necessarily the most desirable outcome, but rather the best that one can do against an opponent of given strength. Figure S.2 shows the optimum strategy for each side of the simple SEAD game. The cross on the right shows that Blue's optimum strategy is to fly strike aircraft 45 percent of the time, while Red's optimum strategy is to fire SAMs only 9 percent of the time. If Red fires SAMs more often, then Blue can move the outcome toward the lower left corner by increasing his SEAD. If Red fires fewer SAMs, then Blue can move toward the upper right corner by flying more strike aircraft. (See pages 6-7.)
- **Determine the expected result of the game.** Having found each side's optimum strategy, analysts check to see whether the outcome of the encounter favors Blue or Red. In the SEAD example, Blue gains an average score of .45 even if Red plays his optimum strategy. This type of finding is very important to analysts because it indicates that even if both sides play intelligently, Blue will have a positive result. Moreover, it indicates that if Red correctly ascertains his situation, then he would rather not participate at all. (See page 7.)

A benefit of game theoretic analysis is that it tells planners how much capability they would need in order to achieve the best outcome for their side. Planners can adjust MoEs to represent better and more valuable aircraft, more costly SAMs, more valuable targets, and other variables. For example, if Blue's strike aircraft were to carry nuclear bombs, then analysts might raise the value of striking the target from 1 to 100. In this scenario, Blue would need to fly strike aircraft only 5 percent of the time and would gain an average overall score of 4.54. Red could do nothing to prevent this loss, even if he fires SAMs 90 percent of the time (his optimum strategy).

This document provides additional illustrations to show how planners can adapt game theoretic analysis to TCT operations such as more complex SEAD operations, interdiction of moving forces, and anti-TBM missions. (See pages 14-39.)

WHEN PLANNING TCT OPERATIONS, ANALYSTS SHOULD CONSIDER THE EXPECTED DURATION OF THE WAR

As described above, planners may use game theoretic analysis to model single TCT engagements with fixed MoEs. However, MoEs may change over the course of a conflict, depending on how long each side expects the war to last. Therefore, when analyzing real-world situations, planners should also consider the implications of longer conflicts with multiple engagements. For example, if a combatant expects a short conflict, then he is likely to place a higher value on achieving his operational objectives, such as striking a target or shooting down an enemy aircraft, than on preserving his own capabilities, such as aircraft and SAMs. This is the case in the nuclear air strike scenario described above. Conversely, if a combatant anticipates a long conflict, then he is likely to place a high value on his limited capabilities and to use them more sparingly until the last few engagements of the war. (See pages 42-49.)

Game theoretic analysis shows that if opponents hold different views about the duration of the war, the side that is correct will have an advantage. Each side will play what he believes to be the optimum strategy and will be surprised to find that his opponent does not play as expected. For example, if Blue correctly anticipates a short war (i.e., Red's targets are high-value) and Red conserves his SAMs, then Blue can inflict heavy damage with a small proportion of strike aircraft. If Blue is incorrect (i.e., Red's targets turn out to be low-value), then Red can shoot down Blue's much-needed strike aircraft with only a small proportion of SAMs. In a real conflict, an intelligent commander faced with this scenario would ask himself whether his opponent is mistaken or whether his opponent knows something that he himself does not. In a planning context, this scenario underscores the importance of understanding MoEs from the opponent's point of view. (See pages 50-52.)

A Simple Game-Theoretic Approach to SEAD and other Time Critical Target Analyses

April 2003

**Tom Hamilton
Richard Mesic**

1 0403 RAND

The effectiveness of Air Force attacks on time critical targets (TCTs) often depends on decisions made by the adversary. For example, air defense radars can be located and attacked much more easily if they are emitting radiation than if they are not emitting and hidden. Vehicles can be more readily found and attacked if they are moving in a column on a road than if they are dispersed and under cover.

This situation presents a challenge to analysts attempting to evaluate the effectiveness of possible technologies and concepts of operations for attacks on TCTs. We must have a way to study likely changes in enemy behavior resulting from various attack strategies and tactics.

One approach to this problem is to use tools derived from game theory. In game theory, a set of mathematical techniques is used to calculate the best strategy to follow in situations in which the outcome of an engagement depends on enemy behavior that we can neither directly control nor reliably predict.

Many military situations can be usefully modeled as two-player, zero-sum games, where zero sum means that the interests of the players are diametrically opposed. The mathematics of this type of game is simple and straightforward. In this documented briefing, we apply well-established methods from game theory to current Air Force problems involving attacks on time critical targets.

Briefing Outline

- ⇒ • Simple game theory
 - Application to data from recent Suppressor runs
 - Application to attacks on moving fielded forces
 - Application to defense against theater ballistic missiles (TBMs)
 - Consideration of multiple engagements

2 0403 RAND

We will first outline the basic mathematics for a straightforward two-player game and show how the mathematics can be applied to a simple examination of tactics for the suppression of enemy air defenses (SEAD).

We then discuss application of this method to a more realistic detailed model of SEAD engagements, to interdiction of moving enemy ground forces, and to defense against theater ballistic missiles.

Finally, we will discuss the complexities that are added to the model by considering how winning strategies evolve during a prolonged conflict.

Example of Simple Game Theory

Consider a two-sided contest in which the participants' interests are diametrically opposed and each side has two possible strategies from which to choose.

For example, consider a simple one-on-one SEAD game. Blue can deploy either a strike aircraft, such as an F-16 CG, or a SEAD aircraft, such as an F-16 CJ. A Red surface-to-air missile (SAM) can either engage or not.

Let us postulate values for a Blue measure of effectiveness (MoE) for different outcomes:

- If Blue sends a strike aircraft and the SAM does not fire, Blue receives a benefit of one.
- If Blue sends a strike aircraft and the SAM engages, Blue receives -5 (the plane is likely to be shot down).
- If Blue sends a SEAD aircraft and the SAM engages, Blue receives 5 (for likely killing the SAM).
- If Blue sends a SEAD aircraft and the SAM does not engage, Blue receives zero.

3 0403 RAND

We model a simple one-on-one tactical engagement. An actual conflict would consist of many engagements. For simplicity, we refer to a conflict between a single aircraft and a single SAM, although the logic can be applied to any simple one-on-one conflict, as between a two-ship flight and a SAM battery.

We postulate a two-sided game in which each side can choose between two possible moves. Each side chooses its move in ignorance of the choice of the other side.

The combination of the decisions of the two sides determines the outcome. Clearly, if either side had the option to make a move that was always best, regardless of the other side's action, that side would always make the move and the game would be uninteresting. We are interested in situations, like the children's game scissors-rock-paper, in which each side can gain an advantage if it can accurately guess the other side's move.

We describe a simple game based on a possible engagement between an aircraft, which may or may not be a specialized SEAD aircraft, and a SAM, which may or may not turn on its radar and engage.

Matrix Describing Simple Air Defense Game		
	Red SAM engages	Red SAM does not engage
Blue flies strike	-5	1
Blue flies SEAD	5	0

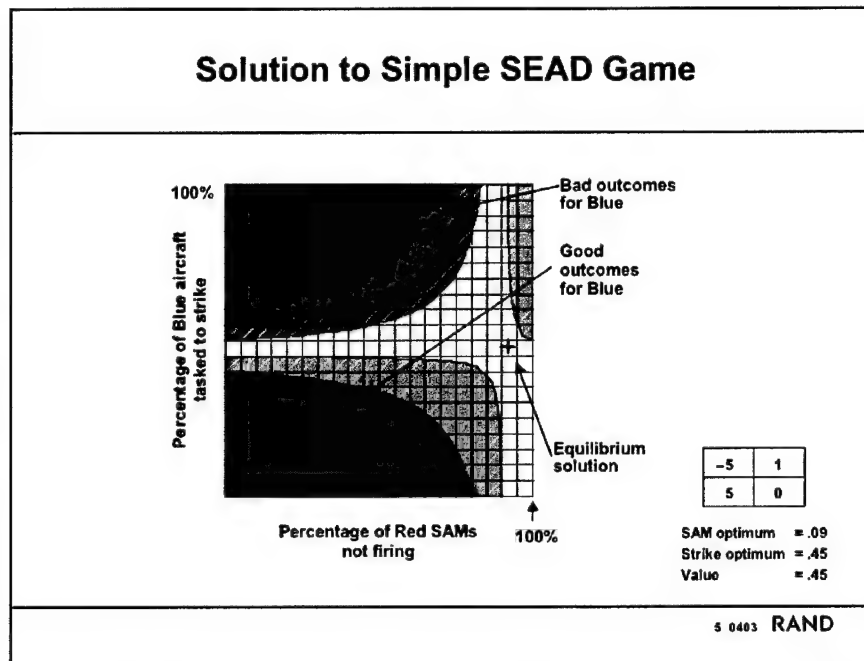
Note: Outcomes reflect value from the Blue point of view. Red is assumed to have *exactly opposite MoEs*.

4 0403 RAND

The above matrix represents the outcomes of the simple SEAD game. The key point is that the outcome represents utility from the point of view of Blue. The utility of Red is assumed to be exactly opposite. If Blue flies SEAD and the Red SAM engages, Blue receives a score of +5. In our assumption, Red's score is -5. That is, Blue's score and Red's score are assumed to sum to zero. This is the meaning of the technical term "zero-sum game."

The property of this matrix that makes it appropriate for game theoretic analysis is that there is no obvious strategy for either side. Both sides could gain an advantage if they had advance knowledge of their opponent's action. If this were not true—if, say, the bottom row of numbers were all positive numbers and the top row were all negative—then Blue would always choose the action associated with the bottom row. In that case, the problem would be simple and the tools presented here inappropriate. If a commander has an option that always produces the best result no matter what the enemy does, the commander should just do it.

The fact that we consider only two-by-two problems is not important. Problems with large numbers of possible actions are perfectly tractable with the techniques presented here. The arithmetic just gets messier.



The above chart contains a great deal of data. Understanding the meaning of this chart is critical to understanding the overall meaning of this briefing, and the balance of this documentation will present data in the format shown here. We will discuss this chart in detail.

For the game described by the matrix on the previous chart, reproduced in the lower right-hand corner of this chart, neither Red nor Blue has one unique correct move. Indeed, consistently making the same move, following what game theorists call a "pure strategy," would allow the opponent an advantage. Consistently making the same move would allow one's opponent to know one's move in advance and exploit that knowledge. In this sort of game, it pays to keep the opposing player guessing.

In the language of game theory, the correct strategies for playing this sort of game are "mixed strategies," in which a player plays a random mix of the available pure strategies. The chart shows the outcome of the game as a function of all possible mixed strategies by Red and Blue. The horizontal axis represents all possible Red mixed strategies. At the extreme left, Red has ordered the Red SAM batteries to fire at every opportunity.

At the extreme right, Red SAMs always hold their fire. Analogously, the vertical axis of the plot represents the Blue strategy. At the top of the plot, Blue has followed a strategy in which Blue aircraft are all tasked to conduct strike missions. At the bottom of the plot, Blue has tasked all Blue aircraft to SEAD.

On the field of the plot, the expected average outcome of the engagement is color-coded. The red area is best for Red; the blue area is best for Blue. Intermediate results are shown by other colors of the spectrum.

The corners of the chart show the outcomes for pure strategies. For example, the upper left corner of the chart corresponds to the case in which Blue always flies only strike aircraft and never flies any SEAD, and Red SAMs always engage. This is a very favorable situation for Red, as shown by the red color on the chart. In contrast, the upper right corner of the chart corresponds to the situation in which Blue only flies strike aircraft and Red SAMs never engage. This is somewhat favorable for Blue, as shown by the green color.

The outcome of the engagement is described by the color at a point (x, y) on the graph. The value of x is chosen by Red and the value of y is chosen by Blue. If the game is repeated several times, both Red and Blue can change their selection in an attempt to improve their outcome. For example, the upper left corner is very unfavorable to Blue, so if it appeared that it was describing the situation, Blue could adjust strategy to put the result in the lower left corner. Blue has full control of the position of the engagement on the vertical axis. But this would create a situation very unfavorable to Red, which Red could improve by changing strategy to move the result to the lower right corner.

Clearly, both players will want to predict the other's strategy. It is likely that one or the other will have a superior capability to do this.

A very important point, which is true although neither intuitive nor obvious, is that either player always has the option of selecting a strategy that gives a guaranteed result, regardless of the opponent's play. This strategy is called a "Nash equilibrium" and is labeled on the chart as the equilibrium solution. On the chart, one can see that if either player has chosen the equilibrium solution (in the yellow area at the right and center of the chart), any change of either x or y will still leave the solution in the yellow area. Unless a player has high confidence that he/she can predict the opponent's moves, the smart move is to play the equilibrium. If a player has reason to suspect that the enemy may have some ability to predict his/her moves, then the equilibrium is definitely the right play. If you play the equilibrium you cannot be outsmarted, even if your opponent knows in advance that you plan to play the equilibrium. We will often refer to the equilibrium strategy as the optimum strategy. The terms are interchangeable in this context.

We will assume here that Red is concerned about the possibility of Blue information superiority. Red will therefore generally play the equilibrium strategy, thereby negating that superiority.

In the particular case plotted, the equilibrium strategy calls for Red to turn on SAM radars 9 percent of the time. This means that the Red battery commander should engage roughly one out of every eleven times an opportunity to engage presents itself. The times that the SAM battery engages should be chosen in a completely random fashion.

For many people, the thought of a military commander acting in a wholly random fashion will seem counter-intuitive, if not bizarre. Yet such actions are recommended by game theoretic analysis. The point is that some specific operational decisions do not matter a great deal in themselves. An air defense commander might decide to order SAMs to shoot on Monday and hide on Tuesday, or to shoot on Tuesday and hide on Monday. In itself, the decision may not make much difference. But one is often at a huge disadvantage if the enemy knows one's plans. If an enemy knows in advance on which days one's SAM will be hiding, that enemy's operations can be made much more effective.

A second issue is that the information available to a commander may not be completely reliable. It is possible that the enemy has deliberately provided a commander with false information in an effort to manipulate the commander into mistaken decisions. This could be a traditional feint or a computer or network attack. Acting randomly provides protection against enemy information warfare.

In our analysis, we assume that both sides appreciate the great advantage of surprising the enemy and either act randomly or have such good security that their actions appear random to the enemy.

Finally, we note that the game has a value, the long-term average of the results of optimal play. In this case it is .45. The fact that this number is positive means that the net advantage is to Blue. If both sides play intelligently, Blue will win. Blue will win only .45, less than half the value (1.0) Blue could get if there were no Red SAMs present, but still a win. Red would rather not play this game at all.

When we discuss Red's optimum strategy, we are not necessarily discussing a winning strategy—just the strategy for obtaining the best outcome available to Red.

Calculation of Optimum Red Air Defense Strategy

The optimal Red strategy in the preceding game is for Red SAMs to fire a fraction of the time equal to

$$\frac{\text{Value of strike}}{\text{Value of strike} + \text{Value of aircraft} \times Pk_S + \text{Value of SAM} \times Pk_A}$$

For the numbers used above, this equals $1/(1+5+5) = 1/11 \approx .09$

This result implies that increases in Blue aversion to casualties and the lethality of SEAD munitions will both tend to drive down the optimum Red air defense engagement rate.

8 0403 RAND

We can formally calculate the optimum (or equilibrium) Red strategy for this simple game. The numerator is the value of the strike—how much damage the Blue aircraft does if it is a strike aircraft not engaged by Red air defense. The denominator is the sum of the value of the strike, the value of an attack against the aircraft, and the value of an attack against the SAM. We consider the value of an attack to be the product of two terms: the value of the target (the loss to the owner if the target is destroyed) and the probability of the target being killed if it is attacked, here written P_k (actually P_{k_S} and P_{k_A} for the P_k of SAMs and anti-radiation missiles, respectively). This means that in our numerical example the numbers refer to the value of the shot. A low P_k shot against a high-value target could have the same value as a high P_k shot against a low-value target.

The major factor driving Red's decision about how often to shoot is the value of the strike Blue may conduct. If the strike is very valuable, Red should shoot. If the value of a strike is lower, in particular if the value of the strike is less than the value of an attack on the SAM or on the attacking aircraft, then Red will engage less often.

It is intuitive that placing a high value on the SAM will make Red more conservative about its employment. It is less intuitive that placing a high value on shooting down a Blue aircraft would drive Red to a conservative strategy. Yet, increasing the value of the aircraft drives Blue to devote more assets to its protection, thus making an actual attack more dangerous for Red. Red benefits by letting Blue devote assets to SEAD that would otherwise be available for strike.

Calculation of Optimum Blue SEAD Strategy

The optimal Blue strategy in the preceding game is for Blue to allocate a fraction of his forces to SEAD equal to:

$$1 - \left[\frac{\text{Value of SAM} \times Pk_A}{\text{Value of SAM} \times Pk_A + \text{Value of aircraft} \times Pk_S + \text{Value of strike}} \right]$$


For the numbers used above, this equals $1 - 5/(5+5+1) = 1 - 5/11 \sim .55$

This result implies that increases in Red SAM effectiveness will cause Blue to allocate more resources to SEAD, and, as the value of the strike increases, the fraction of forces allocated to strike decreases. The second result is highly counter-intuitive.

9 0403 RAND

We also present the explicit formula for the optimum Blue strategy. One result is that as the effectiveness of the SAM system increases, so should resources devoted to SEAD. Another result, which is less obvious, is that as the value of the strike increases, the sorties devoted to SEAD should also increase. This is a response to the fact that as the value of the strike increases, one can expect Red to use SAMs more aggressively, thereby increasing both the need for the protection of SEAD aircraft and the opportunity to kill SAMs.

Cold War (Nuclear) Air Defense Game




10 0403 RAND

We now consider what the game looks like if the value of a strike is extremely high. This would describe a situation in which each aircraft that penetrated enemy air defenses could do an enormous amount of damage, perhaps by using nuclear weapons or other weapons of mass destruction (WMD). We see that the equilibrium solution is near the lower left corner of the graph. Red should use SAMs aggressively, and Blue should respond by providing strikers with a heavy SEAD escort.

We note that the value of this game is 4.54, very favorable to Blue and a direct result of the high effectiveness of Blue's strike.

This is certainly an important case, but it is not the only case. In situations in which Red has some ability to ride out Blue strikes, the optimum strategies are very different.



We next examine a case in which the value of each individual strike is low compared to the value of the aircraft and the value of the SAMs. It can be considered to correspond to a situation such as Operation Allied Force. In this case, Red should fire very rarely—one half of one percent of the time for the values postulated. At the same time, Blue should make a heavy commitment to SEAD. If a Red SAM were able to make an opportunistic attack on an unescorted Blue aircraft, the cost to Blue would be very high.

It is interesting to compare the optimum strategies predicted by game theory for the case of very valuable aircraft and SAMs with the description of the air war over Serbia by Admiral James Ellis, Commander, Joint Task Force Noble Anvil. Admiral Ellis wrote, "Redundant [air defense] systems and well-trained operators with the discipline to wait for a better opportunity affected tactical employment of airpower throughout the war. This required significant intelligence, surveillance, and reconnaissance (ISR) and SEAD efforts throughout the campaign."¹ (This simple game does not factor in other strategy options for either side such as Blue deciding to use more standoff weapons to improve strike effectiveness while minimizing exposure to Red SAMs. Variations of this sort could be analyzed in additional games.)

¹*Air War Over Serbia*, Vol. 1, *Aerospace Power in Operation Allied Force*, Headquarters, United States Air Force, October 2000, p. 433.

Questions About Red and Blue MoEs

What if a side adopts an incorrect strategy but has a correct understanding of the MoEs?

- The side will be punished in a way that it will understand. By trial and error it will converge on the correct strategy.

What if the two sides have different MoEs (i.e., the MoEs are not diametrically opposed)?

- There are two cases:
 - One side may have made a mistake in estimating correct MoEs.
 - That side will be punished in ways it may not understand. This is very dangerous.
 - The two sides may have genuinely non-exclusive interests.
 - This produces a complicated situation in which both sides may benefit from explicit negotiations. Such negotiations can be profitably studied by game theory, but the problem is quite complex.

12 0403 RAND

Let us consider the results if the two sides do not have a perfect understanding of the situation. If one or both sides understand the goals of the campaign but adopt incorrect strategies, they will be punished for their error in a way that they will understand. They will not be getting what they want and will adjust their strategy accordingly. According to the *Gulf War Air Power Survey* (GWAPS), in Operation Desert Storm "after the first Air Tasking Order day, Iraqi SAM operators became more and more hesitant to employ their weapons against Coalition aircraft in anything approaching a normal, let alone optimum, way."²

Note that the GWAPS authors use the word "optimum" in a different fashion that it is used in this documented briefing. They mean firing in such a way as to produce the greatest immediate U.S. losses. As GWAPS notes, "guided firing incrementally declined as the war progressed." The authors of this report would argue that from a game theoretic viewpoint, the Iraqi strategy was, unfortunately, becoming "more optimal" as they developed a better understanding of their overall situation.


A more complicated situation results if the two sides in the conflict have a genuinely different assessment of the value of different results. There are two cases. One side could simply be wrong. It may have misjudged the effects needed to produce victory, which is likely to lead to disaster. One can argue that the U.S. involvement in Vietnam, especially in

² Eliot Cohen (ed.), *Gulf War Air Power Survey*, Office of the Secretary of the Air Force, Washington, D.C., 1993, Vol. II, Pt. II, p. 141. See also chart 3-9 on p. 138.

1965–1968, suffered from this type of error. The United States believed that killing Vietcong would produce victory. North Vietnam did not place as high a value on the lives of individual Vietcong as the United States believed it did. The United States adopted incorrect strategies as a result.³

Another situation is if the goals of the two sides are not genuinely exclusive. In this case, both sides can benefit from a negotiated settlement. This is a very important branch of game theory, but we do not consider it here. This report addresses the problems of airmen, not diplomats.

³ R. S. McNamara, J. G. Blight, and R. K. Brigham, *Argument Without End*, Public Affairs, New York, 1999, p. 216.

Briefing Outline	
	• Simple game theory
	• Application to data from recent Suppressor runs
	• Application to attacks on moving fielded forces
	• Application to defense against TBMs
	• Consideration of multiple engagements
14 0403 RAND	

Having outlined a simple game theory approach to understanding the SEAD problem, we now apply the approach to detailed, engagement-level modeling.

Simple SEAD Game

We consider a two-player SEAD game whose payoff matrix is:

	Red SAM fires	No Red fire
Blue striker	-a	b
Blue SEAD	c	0

For such a game, the optimum Red strategy is $\frac{b}{a+b+c}$

For such a game, the optimum Blue strategy is $\frac{c}{a+b+c}$

The value of this game to Blue is $\frac{cb}{a+b+c}$

15 0403 RAND

With this chart we review the basic results of game theory.

Application to Cases Run in Suppressor

Suppressor runs at RAND have sometimes considered Blue attacks in two waves, a destruction of enemy defenses (DEAD) wave followed by a DEAD + Strike wave.

The DEAD wave consists of two phases. In phase one, Blue decoys excite emissions from Red SAM radars. In phase two, Blue detects the locations of the emitting radars and attacks them with cruise missiles.

The DEAD + Strike wave also includes the decoys and subsequent missile attack, but the decoys are accompanied by strike fighters that attack Red targets.

We assign numbers to the value of each entity potentially lost and the number lost in each wave:

	Value	Lost in DEAD	Lost in DEAD + Strike
Blue cruise missiles	1	100	100
Blue aircraft	1000	0	.5
Red SAMs	100	3	3
Red targets	10	0	3

If Red chooses not to engage, no aircraft or SAMs are lost. Cruise missiles are still used; three Red targets are still lost.

15 0403 RAND

We now apply the apparatus of simple game theory to the results of some engagement-level modeling. The point here is not the precise scenario modeled, just its general characteristics and results.


We model an attack against targets in an area defended by a modern air defense system, including SA-10-equivalent SAMs. In our model, Blue has two options. Blue can send unmanned decoys to overfly Red SAMs and follow up with cruise missile attacks on any Red radars that emit. The decoys are followed by four strike aircraft that hit targets in the area Red is attempting to defend. Or Blue can just send the decoys and cruise missiles, putting no pilots at risk and not striking any Blue targets other than the emitting SAM radars. Red has the choice of engaging with SAMs or simply ignoring Blue's actions. We have assigned values to Blue's cruise missiles and aircraft and to the Red SAMs and targets. Note that while the engagement-level model can compute weapons effectiveness, Pk, the assignment of value to a military asset is a more subtle decision. Although game theory requires us to make such assignments of value explicitly, in practice they are often not made explicitly. Indeed, it is impossible to precisely quantify the value of, say, a pilot's life compared with the value of a bridge on the target list. Nevertheless, such uncomfortable human-value judgments, even if never spoken out loud, form the underpinning of all military operations.

Values in Game Matrix for Suppressor Case		
	Red engages	Red does not engage
Blue sends DEAD wave	-100 CM 300 SAMs	-100 CM
Blue sends DEAD + Strike wave	-100 CM 300 SAMs -500 aircraft lost 30 targets	-100 CM 30 targets
CM = cruise missile		
17 0403 RAND		

Based on the values we assigned in the previous chart and outputs from an engagement-level model, we can fill in a game matrix describing the options of both sides and the value of the outcomes.

If Blue sends in only decoys and SEAD cruise missiles and Red engages, Red will lose three SAM batteries. If Blue sends in only decoys and SEAD cruise missiles and Red does not engage, the only result is the cost to Blue of the cruise missiles (100 in the units used here). If Blue also sends in a strike package and Red engages, Blue has a 50 percent chance of losing an airplane, adding a negative 500 for Blue's score. But if Blue sends in strikers, three targets are struck, adding a positive 30.

Obviously, the matrix numbers would be different if the value assignments were different. The great virtue of this approach is that it forces the analyst to address the value question directly.



This chart presents the results for the game described on the previous charts. When one adds up the values for each matrix element, one computes the matrix shown here. The solution is shown on the right in familiar form. Red's optimum firing rate is very low, 6 percent. Most important for Blue, Blue's expected game value is low, -82. The problem is that Blue's concept of operations (CONOP) relies on the heavy use of expensive cruise missiles for SEAD. Red's strategy is to usually ignore Blue's strike aircraft and just hope that Blue wastes expensive resources unsuccessfully chasing SAMs that cannot be killed because they are not engaging.

Ultimately, this is a bad game for Blue. If this is the best CONOP Blue can devise, Blue should not attack. On the other hand, if the input values were different, this CONOP could be attractive for Blue. For example, significantly decreasing the cost of the Blue cruise missiles or greatly increasing the value of the strikes would give the game a positive value for Blue.

Mathematical note: Because this game matrix has a non-zero value in the lower right corner, the arithmetic is a little more complex than the formulas presented elsewhere in this briefing.

When Should Blue Play?

The game becomes of neutral value if the value of the cruise missile is reduced from 1 to .18. This does not change the strategy for either side. (The cost of the cruise missiles is included in all four outcomes, so it has no effect on the optimum decision.) Even if Blue had an unlimited supply of cost-free cruise missiles, the correct Red strategy would still be to fire 6 percent of the time.

19 0403 RAND

Blue could make the game more attractive if Blue could reduce the cost of the cruise missiles. If the cost of Blue's SEAD cruise missiles could be reduced to 18 percent of the baseline cost, the value of the game would be neutral. If Blue's cost could be reduced further, Blue would come out ahead. But even if the cruise missiles were free, Red's optimum strategy would not change. Red should still refuse to engage 94 percent of the time.

Briefing Outline

- Simple game theory
- Application to data from recent Suppressor runs
- ⇒ • Application to attacks on moving fielded forces
- Application to defense against TBMs
- Consideration of multiple engagements

20 0403 RAND

In this section, we consider application of simple game theory to attacks on fielded forces such as tanks, trucks, and other military vehicles moving on the ground.

We will consider that enemy ground forces have three possible courses of action. They can remain stationary and hidden, they can move across country in dispersed formation, or they can move along roads. The more quickly they move, the more exposed they are to aerial attack.

Two CONOPs for Simple Spreadsheet Analysis

Three systems

- Global surveillance system (Global Hawk, satellite, etc.)
- UAV with target identification capability
- Shooter (fighter or UCAV)

Two UAV CONOPs

- Respond to cues, identify targets
- Patrol, identify targets

Two integrated CONOPs

- GSS - UAV - Shooter
- UAV patrol - Shooter

UAV = unmanned aerial vehicle

UCAV = unmanned combat aerial vehicle

GSS = Global Surveillance System

21 0403 RAND

We consider two Blue CONOPs for attacks against Red vehicles. We use a simple spreadsheet model to evaluate the effectiveness of the two CONOPs.

In the first CONOP, some sort of system capable of surveying a large area continuously observes the Red area of operations. This system could be any system capable of large area surveillance. When the system detects a potential target, it passes the target position to a UAV, which flies to the target position and investigates the potential target to determine if it is appropriate for attack. If it is, a fighter is vectored to the target location and attacks. This CONOP is potentially effective against any moving Red vehicle in the area of operations.

Alternatively, the UAV can operate autonomously, searching for targets on its own and calling in the fighter as appropriate. This CONOP is only effective against Red vehicles moving along known lines of communication. We have implicitly assumed that the UAV lacks the capability to efficiently observe large areas, and that the wide area surveillance system lacks the ability to reliably identify targets.

Model Parameters	
Terrain <ul style="list-style-type: none"> • Battlefield dimension D x D kilometers Terrain exposure and hides <ul style="list-style-type: none"> • Mean length of exposure, T_e Rate of enemy target exposures <ul style="list-style-type: none"> • Exposures per day, E_d Target concentration <ul style="list-style-type: none"> • Percentage of targets in the 1 percent of the total area that is most crowded False target percentage Global surveillance <ul style="list-style-type: none"> • P_{det} per target per hour (tracked till hidden) 	UAV speed UAV field of regard UAV time to identify Loss rate for unobserved targets (if UAV leaves) (percent/hour) UAV percent false positive UAV percent false negative Fighter velocity Fighter field of regard Fighter time to identify Fighter percent false positive Fighter percent false negative Fighter probability of kill
22 0403 RAND	


The chart above lists the parameters that are input into our simple spreadsheet model. The details of the model are not important to this argument. We are simply examining the broad, qualitative nature of the results.

Measures of Effectiveness	
Enemy targets killed per hour	
Potential targets detected per hour	
Percent exposed enemy killed	
Percent exposed enemy detected	
False targets killed per hour	
Number targets/hour investigated by UAV	
Number targets/hour investigated by fighter	
23 0403 RAND	

Our model calculates several measures of effectiveness. The fundamental one is enemy targets killed per hour. We measure this as a function of both Red activity (for example, vehicle exposures per hour) and Blue CONOPs, either surveying the battlefield with a theater surveillance asset or patrolling a known line of communication.

Battlefield Geometry

Blue UAV may be patrolling the LoC or responding to cues



Red forces may be moving along LoC or cross-country

We model a diamond-shaped area with a line of communication running through it. The width of the diamond is two-thirds its length.

Red units may move quickly down the LoC, in which case they are vulnerable to identification by UAV patrols or UAVs responding to cues. They may also move cross-country, in which case they move more slowly but are only vulnerable to cued UAVs.

LoC = Line of communication

24 0403 RAND

This chart summarizes the geometry of the situation modeled. If Red vehicles are moving cross-country, they move slowly and can be found only by UAVs receiving support from global surveillance assets. Alternatively, the Red vehicles can move down the line of communication (LoC). In this case, they move more rapidly but are vulnerable to being detected by both global assets and UAVs patrolling the LoC.

This outlines the structure of the two-player, two-strategies game we are studying. Blue can use global assets or simply patrol the known LoC. Red can move rapidly along the LoC or more slowly cross-country.

Cued Versus Patrolled CONOPs

Parameter values:

Target concentration = .2 (proportion along LoC)

UAV field of view (FOV) = 1 km²

UAV speed = 100 km/hr

Fighter speed = 600 km/hr

Cue duration = 1 hr

Fighter Pk = 1


Global Pdet = .1 /target/hour

Pdet = Probability of detection

25 0403 RAND

We define some typical parameters we need to derive a quantitative solution. These parameters are notional and are not meant to describe any specific systems.

Simple Spreadsheet Outputs



Results against 20% of Red force moving on LoCs

26 0403 RAND

The plots in the chart above show the result of the simple spreadsheet model for the two alternative CONOPs as a function of the rate of Red exposure. We have assumed here that 20 percent of Red vehicles that move, and are therefore exposed to potential detection, are moving along LoCs known to Blue.

Note that the target kills per hour is better for the global surveillance CONOP when the enemy activity is at a relatively low level. However, when the enemy becomes very active, patrolling LoCs becomes more efficient. There is no advantage for a shooter to search for targets in the countryside when the road is jammed with enemy vehicles.

Comparison of Patrol and Cued CONOPs

For the default parameters examined above:

Cued operation produces 1 kill for each 40 exposures

Patrol operation produces 1 kill for each 10 exposures

Cued operation is effective against all moving enemy forces, on LoC or dispersed

Patrol operation is effective only against enemy moving along LoC

Consider a case in which a Red force of 40 vehicles wishes to move through the area under interdiction. It can move either along the LoC or in a dispersed pattern. We assume that moving a vehicle creates one exposure per hour. Let us assign each vehicle a value of 10. We further assume that moving each vehicle along the road or in a dispersed pattern for an hour produces values to Red of .25 and .1, respectively.

The game matrix is:

	Red dispersed	Red on LoC
Blue patrols	$0 - 4 = -4$	$40 - 10 = 30$
Blue cued	$10 - 4 = 6$	$10 - 10 = 0$

27 0403 RAND


Just as we earlier used results from the Suppressor SEAD model to generate input values for a game theoretic analysis, we can use results from the simple interdiction model described above to generate inputs for a game theoretic analysis of the interdiction problem.

We can use the results to quantify the relative effectiveness of cued operations versus patrol operations in a situation in which Red can vary CONOPs in response to Blue's actions.

We assume that moving a vehicle along a road provides a value to Red of .25 and that a Red vehicle gains a value of .1 by moving cross-country. We assign each Red vehicle a value of 10 and in this way construct a game matrix similar to the matrices constructed earlier.

As always, positive numbers are good from a Blue point of view and Red's point of view is assumed to be exactly opposite.

Game Solution



Red optimum = 75% dispersed, 25% LoC
Blue optimum = 85% cued, 15% patrol
Game value = 4.5 (for Blue)

Important point is that game value is positive for Blue. Red should not play at all.

This is victory. Blue has created a situation in which the correct Red strategy is to do nothing. This is all that airpower can deliver.

28 0403 RAND

We solve the game defined by the matrix in the previous chart and construct the familiar solution space. We find that Red benefits from following a strategy in which 75 percent of Red vehicles move cross-country. Blue's optimum strategy is to use 85 percent of available assets to support operations cued by theater surveillance assets. In effect, a few Blue assets patrolling the LoC will drive most Red vehicles into the country.

The net value of this game is positive for Blue. Red would do better not to play. Red should keep its assets under cover. In effect, Red should halt all offensive operations.


Nature of Blue Victory

Game Matrix	Red dispersed	Red LoC
Blue patrols	$0 - 4 = -4$	$40 - 10 = 30$
Blue cued	$10 - 4 = 6$	$10 - 10 = 0$

Increasing Red reward for moving will eliminate Blue advantage

Doubling Red's value for moving to .5 per hour on LoC and .2 dispersed yields:

Game Matrix	Red dispersed	Red LoC
Blue patrols	$0 - 8 = -8$	$40 - 20 = 20$
Blue cued	$10 - 8 = 2$	$10 - 20 = -10$



Red optimum = .25 LoC
 Blue optimum = .3 patrol
 Game value = -1

Red will play if Red considers that moving tanks on the LoC or dispersed is worth 5% or 2% per hour attrition, respectively.

29 0403 RAND

Consider the result from the previous chart—that in the situation we have postulated, Red should cease offensive operations.

What would have to change for Red's correct strategy to be to move its vehicles from their hides and resume forward movement? The key is to increase the utility that Red derives from movement relative to the disutility Red derives from losing vehicles.

We can compute that Red will continue to advance if moving vehicles along the LoC or dispersed is worth accepting 5 percent or 2 percent attrition per hour, respectively. The relative value of moving on roads and cross-country has not changed. Red should still move cross-country. The point is that Red should move only if it is very tolerant of losses.

TCT Victory

The best that an interdiction campaign can hope to achieve is total paralysis of the enemy. This is an *effect*.

An appropriate MoE for an interdiction campaign is: "How willing to die must the enemy be before moving becomes a correct strategy?"

Since hiding will work for the foreseeable future, and will often be the correct strategy for an enemy confronted with an effective TCT capability, enemy losses may not be a good measure of TCT capability. Effects-based analysis based on game theory provides a quantitative evaluation of the alternative TCT strategies.

30 0403 RAND

This chart outlines conclusions that follow from the previous analysis. The effect of an interdiction capability is to force Red to consider what sacrifice it is willing to accept in order to continue forward movement. If Blue has a strong position, Red will elect not to move. This is Blue victory.

While Blue victory can be considered to result directly from Red's perception of Blue's effectiveness, it does not mean the perception is any substitute for reality. In this model, the way to create a perception of effectiveness is to be genuinely effective. It also means that if Red foolishly fails to perceive Blue's power, Red will be destroyed directly.

Game theory provides a mechanism for starting with the physics-based capabilities of weapons systems and calculating effects on enemy behavior.

Briefing Outline

- Simple game theory
- Application to data from recent Suppressor runs
- Application to attacks on moving fielded forces
- ⇒ • Application to defense against TBMs
- Consideration of multiple engagements

31 0403 RAND

We now consider an application of simple game theory to the problem of defense against enemy theater ballistic missiles.

Theater Missile Defense (TMD) Game

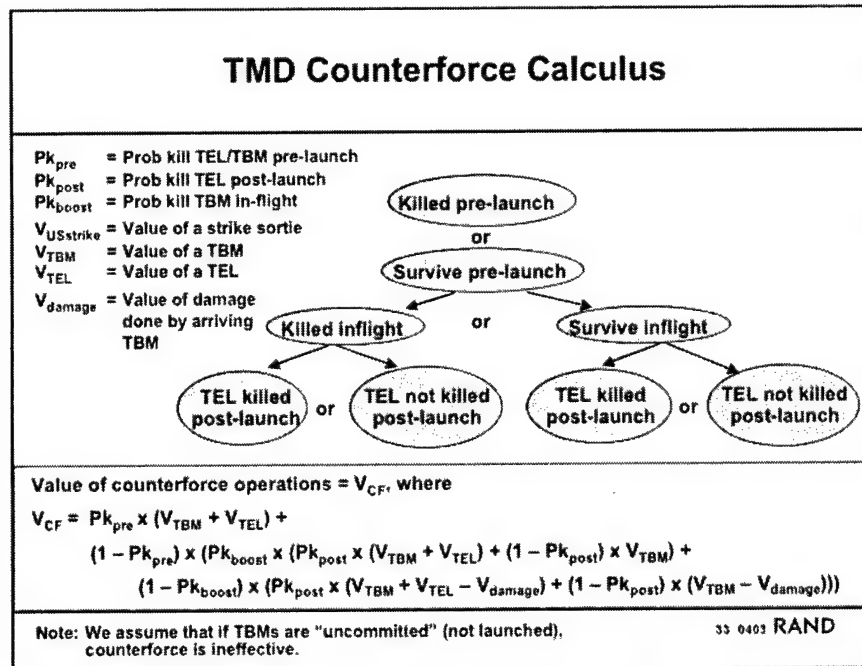
Consider a simple engagement-level TMD game. Blue can allocate strike aircraft, such as an F-15E, either to strike or TMD counterforce missions.

Let us postulate values for a Blue MoE for different outcomes:

- If Blue sends a strike aircraft and the TBMs are not launched, Blue receives a benefit of one.
- If Blue sends a strike aircraft and the TBMs are launched, Blue loses his target value at risk (dependent on TBM probability of defense penetration and damage potential).
- If Blue sends a counterforce aircraft and the TBM launches, Blue receives value dependent on the effectiveness of his counterforce systems (see next chart).
- If Blue sends a counterforce aircraft and the TBM does not launch, Blue receives zero.

32 0403 RAND

We develop values for possible outcomes of an engagement in which Red may either fire a TBM or not and Blue may either commit an aircraft to anti-TBM counterforce or commit that same aircraft to some other valuable mission.



33 0403 RAND

We can fill in the values of the game matrix by making plausible assumptions about the values of the input variables. The choice of values for the assets placed at risk may seem arbitrary. It is important to note that any decisionmaker making these sorts of decisions must implicitly place values on possible outcomes. For example, we have assumed that the damage done by a successful Red transporter-erector-launcher (TEL) attack is 10 times the value of the damage done by a successful Blue fighter attack on a non-TEL target. It is easy to imagine that this number could be greater or lesser depending on the political situation. Certainly, if the TEL were armed with WMD, the input numbers would be different.

As the chart shows, we have included the probability that a launched TBM can be killed inflight in the defense-in-depth game. By "inflight," we mean in the boost phase, in midcourse, or in the terminal phase. This integrated active defense parameter is included in our counterforce analyses because it can affect the optimal allocations of air assets to counterforce or to strike missions.

Matrix Describing Example TMD Game		
	Red launches TBMs	Red does not launch TBMs
Blue flies strike	-4	1
Blue flies counterforce	2.25	0

Parameters:
 $P_{k_{pre}} = 0.5$
 $P_{k_{post}} = 0.5$
 $P_{k_{boost}} = 0.5$
 $V_{US_{strike}} = 1$
 $V_{TBM} = 1$
 $V_{TEL} = 5$
 $V_{damage} = 10$

Note: Outcomes reflect value from the Blue point of view.
Red is assumed to have *exactly opposite MoEs*.


34 0403 RAND

The results of this game can be described in the familiar two-by-two matrix that shows the expected value (to Blue) of each of the four strategy pairs. We can fill in the values of this game matrix by making plausible assumptions about the values of the input variables in the equation in the previous chart. We have assumed that the damage done by a successful Red TEL attack is 10 times the value of the damage done by a successful Blue fighter attack on a strike (non-TEL) target. It is easy to imagine that this number could be greater or lesser depending on the political situation. Certainly, if the TEL were armed with WMD, the input numbers would be different.

It can be argued that information dominance might enable Blue to know in advance whether Red will or will not launch a TBM. If such information were available to Blue, the situation would be simplified. Indeed, information dominance may well allow Blue to know the location of planned Red TBM launches before they occur. This would make pre-launch attacks relatively easy. However, Red could defeat such Blue information dominance by a dispersal of decisionmaking power to lower commands who could make tactical decisions on short notice in an uncoordinated way. This would create the more difficult situation modelled here.

Note: We have not in this analysis considered the use of WMD, which would introduce issues reminiscent of the Cold War (stability, offense/defense trades, deterrence, etc.). There is a voluminous literature devoted to those issues.

Solution to Example TMD Game



35 0403 RAND

The area in the chart represents all possible combinations of Blue and Red mixed strategies. The colors distinguish the value contours: red is bad for Blue, blue is good for Blue (and, hence, bad for Red), yellow is in between. As one would expect, the corners of the area have the most extreme values. If Red always fires and Blue always flies counterforce, the outcome is in the lower left corner of the chart and is very favorable to Blue. Alternatively, if Blue never flies counterforce and Red always launches TBMs, the outcome is in the upper left corner and is very favorable to Red.

For the numbers we have shown here, the equilibrium Red strategy is to launch TBMs 14 percent of the time. Blue should fly counterforce 69 percent of the time. The outcome of this game is a value to Blue of .31. This is a positive number, so Blue is winning. But Red is still doing better than if it ceased TBM activity altogether. In that case, Blue could always strike and the value of the game to Blue would be 1. Red is not adopting a winning strategy here. Red is adopting a palliative strategy to minimize loss while prolonging the conflict.

Calculation of Optimum Red Strategy

The optimal Red strategy in the TBM game is for Red TBMs to fire a fraction of the time equal to:

$$\frac{V_{\text{UStrike}}}{V_{\text{UStrike}} + [V_{\text{damage}} \times (1 - Pk_{\text{boost}}) - V_{\text{TBM}}] + V_{\text{CF}}}$$

For the numbers used above, this equals $1/(1 + 4 + 2.25) = 1/7.25 \sim .14$

• Of course, this may result in an all new ball game.

36 0403 RAND

The above result implies that increases in Blue counterforce capabilities will tend to drive the optimum Red launch rate down, but boost-phase (and by implication other active and passive defense) improvements will drive it up. Also, as V_{damage} increases (e.g., via WMD), launch rates should decrease.

Calculation of Optimum Blue Strategy


The optimal Blue strategy in the preceding game is for Blue to allocate a fraction of its fungible forces to counterforce equal to:

$$1 - \left[\frac{V_{CF}}{V_{CF} + V_{UStrike} + [V_{damage} \times (1 - Pk_{boost}) - V_{TBM}]} \right]$$

For the numbers used above, this equals $1 - 2.25/(2.25 + 1 + 4)$
 $= 1 - 2.25/7.25 \sim .69$

37 0403 RAND

This result implies, reasonably enough, that increases in Red TBM threat effectiveness (viz., allied value at risk) will cause Blue to allocate more flexible resources to counterforce, but, as the effectiveness of Blue active/passive defenses increases, the fraction of Blue forces allocated to counterforce decreases.



In this chart, we compare five games corresponding to different Blue TMD architectures. The counterforce capabilities are described by the pair labeled $CF(Pk_{pre}, Pk_{post})$. In the upper left-hand graph, these numbers are $CF(.5, .5)$. The defense number is the overall probability ($Pk_{inflight}$) (0 in the upper left-hand graph) of inflight kill.

Our objective in this chart is not to determine what may or may not be achievable in the real world—such an assessment is beyond the scope of this study. However, it is useful to assess the improvement in outcome that might be achievable with a combination of modest improvements in the individual elements of a theater missile “defense in depth.” To begin, suppose that we have no significant ability to suppress the launch of enemy TBMs but can kill 50 percent of enemy TELs post-launch. Suppose also that we develop and field a missile defense system capable of destroying 50 percent of missiles in flight. The game that would result from this combination of capabilities is as shown in the lower left-hand chart.

Clearly, the resulting capability is not good. In this game, each time Red launches missiles it can be certain of a bad outcome for Blue. Hence, Red should launch missiles all of the time and it makes no difference whether or not Blue chooses to execute counterforce attacks; the outcome is always bad. If for some reason, however, Red does not launch a missile attack, Blue is better off executing strike operations rather than counterforce operations.

What if Blue were able to achieve some significant pre-launch suppression capability? Such a capability might be possible, for example, in situations

where Blue can suppress the launch of enemy missiles by degrading the physical infrastructure used to support ballistic missile operations. The upper left-hand plot describes the results of a game in which Blue is able to suppress (or kill) 50 percent of the missile launches and is also able to kill 50 percent of the TELs after launch.

In this game (upper left), Blue has no defense against missiles in flight. The result is bad for Blue, although somewhat less bad than the previous game. In this game, Red achieves the best result if it launches and Blue does not conduct counterforce operations. When Red launches, Blue achieves an improved (but still bad) outcome if it has chosen to conduct counterforce operations. However, if Red chooses not to launch, Blue achieves a better outcome if it has not chosen to conduct counterforce operations and uses its aircraft to conduct strike operations instead. In this game, therefore, Red is always better off if it launches. If Red launches, Blue is better off if it conducts counterforce operations.

In the real world, equipment failures and other sources of uncertainty might prevent Red from completing a launch every time that it chooses to do so. To the extent that such failures occur, Blue might be able to improve its overall outcome by engaging in a mixed strategy of counterforce and strike operations with its aircraft.

The two games shown in the center graphs describe the effect that very good inflight defenses or counterforce capabilities have on game outcomes. In the center upper plot, Blue has no inflight defense but has excellent pre- and post-launch counterforce capabilities. In the center lower plot, Blue has excellent inflight defenses, no pre-launch counterforce capability, and moderate post-launch capability. Red and Blue should both adopt mixed strategies for these two games, with the expected results for Blue considerably more favorable than in the two games on the left.

In the game shown on the right, Blue has a defense against inflight missiles that has an effectiveness of 50 percent. The result of this game is much better for Blue. In this game, Blue can achieve desirable outcomes if it flies counterforce operations when Red launches missiles or if it flies strike missions when Red does not. The only bad outcome for Blue occurs when Red launches missiles and Blue flies strike operations. The remaining possibility (Red does not launch, Blue flies counterforce) is neutral. Both Red and Blue would use a mixed strategy in this game.

These results suggest that it may be as good to develop competent (if imperfect) balanced defenses in-depth as to attempt to achieve exquisite capabilities in some areas (e.g., boost-phase defense) at the expense of others (e.g., counterforce).

Briefing Outline

- Simple game theory
- Application to data from recent Suppressor runs
- Application to attacks on moving fielded forces
- Application to defense against TBMs
- ⇒ • Consideration of multiple engagements

40 0403 RAND

We now consider the nature of conflicts consisting of multiple, serial engagements. For this section of the document, the mathematics becomes a little more involved, although nothing that we present here is in any way new or controversial. We continue to use well-established methods.

Bibliographic notes: The basic game theory used in this document was developed in the 1920s by John Von Neumann and is authoritatively presented, along with much elaboration, in J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, New Jersey, 1944. Non-specialist readers can find accessible, accurate expositions of the theory in J. D. Williams, *The Compleat Strategist*, McGraw-Hill, New York, 1954, and in M. E. Dresher, *Games of Strategy: Theory and Applications*, Prentice-Hall, Englewood Cliffs, N.J., 1961. Both of these books are in print in Dover editions.

The dynamic programming is a little more modern, developed in the 1950s and summarized in R. Bellman, *Dynamic Programming*, Princeton University Press, Princeton, N.J., 1957.

We note that after the early success of game theory in analyzing military problems, which are typically zero-sum, there was a great expansion in the application of game theory to problems in economics and political science, problems that are usually not zero-sum. This document maintains a strict focus on military operations and does not consider post-1950s intellectual developments.

Simple SEAD Game

We consider a two-player SEAD game whose payoff matrix is:

	Red SAM fires	No Red fire
Blue striker	-a	b
Blue SEAD	c	0

For such a game, the optimum Red strategy is $\frac{b}{a+b+c}$

For such a game, the optimum Blue strategy is $\frac{c}{a+b+c}$

The value of this game to Blue is $\frac{cb}{a+b+c}$

41 0403 RAND

We review the simple SEAD game we discussed at the beginning of this document.

Multi-Engagement SEAD Game

Multi-engagement SEAD games can be divided into two categories based on the relationship of the optimum Red attrition rate, the size of the Red force, and the duration of the conflict.

Optimum Red attrition [ORA] = $\frac{ab + b^2}{(a + b + c)^2} PkH$, where PkH is the Pk for attacks on SAMs.

If $N_{\text{engagements}} \times \text{ORA} < \text{Initial Number Red SAMs}$:

Red and Blue optimum tactics are constant throughout the conflict.

Red can use with maximum efficiency only a number of SAMs equal to $N_{\text{engagements}} \times \text{ORA}$.

Red can use more SAMs, but only with reduced efficiency, i.e., an inferior exchange rate.

Possible procurement implication for Red: increasing SAM quantity should face diminishing returns; increasing SAM quality always wins. (More SAMs are useful, but they can be used only in a suboptimal way.)

42 0403 RAND

Let us consider the effect of multiple engagements upon the simple SEAD game we described on the previous chart and in the first section of this document. Throughout the discussion that follows, we assume that Blue has sufficient resources to replace any destroyed assets and that Red does not. That is, we will focus on the effect on operations of the potential destruction of Red's combat capability. This simplifies the math and reflects the reality of the conflicts in which the United States is likely to be engaged in the near future. U.S. casualties are likely to be a serious consideration in future conflicts.

We start by considering the result of an engagement in which both sides follow the optimum strategy. Following the optimum strategy produces a certain level of Red loss. The question is whether Red can sustain the optimum level of loss throughout the conflict.

The concept of optimum Red loss is not obvious. For any set of value judgments about the value of military resources, there are, from game theory, optimal strategies for both Red and Blue. If both sides play optimally, there is a well-defined loss that Red can expect. This is what we refer to as the "optimum Red loss." Obviously, Red would prefer to have no loss at all, so the optimum Red loss is not Red's favorite possible outcome. It is the loss Red can expect when following Red's optimum strategy.

The question is whether Red can sustain this optimum level of loss throughout the conflict. If Red can do so, Red is in a favorable position and its strategy is clear.

The implication for Red's peacetime procurement is that there is a limit to the quantity of air defense assets a nation can employ with optimal efficiency. More SAMs are always better, but there is a law of diminishing returns. By contrast, better SAMs are always useful. There is no law of diminishing returns with respect to SAM quality.

Of course, this assumes that Red has already procured a force large enough to sustain optimal attrition through a future conflict. Moreover, if the cost of the better SAMs is too high, their cost-effectiveness may be less than that of even a non-optimal buy of less-effective SAMs. Real world situations will be complex.

Lanchester Combat Laws

Lanchester square law: Value of quality is proportional to value of quantity squared

- Appropriate for aimed fire (classical air and naval combat)
- Provided great advantage to superior force

Lanchester linear law: Value of quality is proportional to value of quantity

- Appropriate for unaimed fire or shoulder-space-constrained warfare
 - For example: World War I artillery or Thermopylae
- Favorable to weaker side

Lanchester constant: Quality is everything

- Appropriate for linear law situations when one or both sides are averse to casualties, e.g., guerrilla war, terrorism
- Favored by weak side with attrition-averse opponent

If Red can sustain optimal attrition, Lanchester constant situation applies.
Neither side can bring a larger force to bear effectively.

44 0403 RAND

This chart is a theoretical excursion of the implications of the concept of optimum Red attrition.

Traditionally, analysts have considered two general classes of combat laws. The Lanchester square law is the classical description of aimed fire.⁴ Originally developed for naval warfare, in which ships aim their fire at a specific enemy ship, it has been often used as a description of air-to-air combat. Pilots do not fire air-to-air weapons unless they have a particular target in their sights. In a Lanchester square conflict, the value of weapon quality is proportional to the value of weapon quantity squared.

The Lanchester linear law is the classic description of unaimed fire, such as World War I artillery. The same mathematical expression also describes combat in which opposing forces are limited by shoulder-space constraints, Thermopylae being the classic example. In a Lanchester linear conflict, the value of weapon quality is proportional to the value of weapon quantity.

The concept of optimum Red attrition introduces a new combat law, which we here call "Lanchester constant." This law applies in asymmetrical situations in which one side is extremely averse to

⁴The clearest exposition of Lanchester's argument remains F. W. Lanchester, *Aircraft in Warfare: The Dawn of the Fourth Arm*, Constable, London, 1916. Good modern discussions of Lanchester's formulation can be found in J. Bracken, M. Kress, and R. E. Rosenthal (eds.), *Warfare Modeling*, Wiley, New York, 1995.

attrition. We have developed this concept as a description of some SEAD campaigns, but the same mathematics can describe a terrorist campaign. The point is that by following an optimal strategy, the weaker power can construct a situation in which neither power can profitably bring to bear additional forces. The weaker power will control its own attrition, limiting it to the optimum.

Multi-Engagement with Limited SAMs

Consider the case in which Red does not have enough SAMs to sustain optimal attrition for the duration of the campaign. In this case, the cost of losing a SAM is increased by the absence of that SAM from future engagements. (Note that this effect is significant only if Blue knows that Red has an insufficient SAM force. This is not an unimportant condition. If Red's optimal firing rate is low, Blue may not be able to detect whether Red's force is suboptimal.)

Mathematically, this is solved by first solving for the last engagement of the war and then iterating forward (i.e., dynamic programming with a finite horizon).

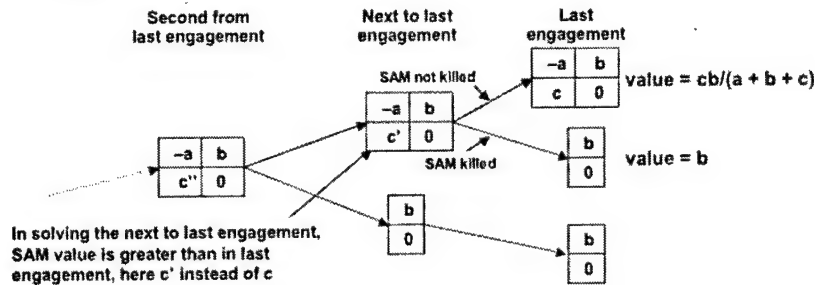
46 0403 RAND

We now consider the alternative—the case in which Red does not have enough SAMs to sustain optimum attrition throughout the conflict. In this case, Red has to consider the opportunity cost of losing a SAM that could be useful in future engagements. This is a well-studied problem that can be solved by dynamic programming.

Dynamic programming is a procedure for determining optimal strategies in a situation that changes over time, partly as a result of strategies adopted previously. It is best described as recursive optimization. Like game theory, dynamic programming was originally developed to address military problems and has gone on to have wide application in economics and industrial organization.

Solving for Multiple Engagements

We consider influence of the result of one engagement on subsequent engagements by constructing a tree. In this case, we consider engagements between one SAM and one replaceable aircraft. If the SAM is killed, it is not present in subsequent engagements, shifting the value of those engagements in Blue's favor. So the value of the SAM is higher for earlier engagements, changing the optimal strategies.




47 0403 RAND

The essence of dynamic programming is to start by considering the situation at the end of the conflict. The last move of a conflict is relatively straightforward. We use the solution we derived earlier to describe this last move. We then consider the nature of the next to last engagement of the conflict. In this engagement, the value of the SAM is increased by an amount that reflects the opportunity cost to Red of not having the SAM available at the last engagement. This is easy enough to compute explicitly. We then turn to the second from the last engagement and compute the value of the SAM at that time. We can follow this process iteratively for an arbitrarily large number of engagements.

The main points are that the value of the SAM increases as we move further away from the last engagement of the conflict and that we can compute this value explicitly.

Change in Optimal Strategies Over Time



We plot the change in Blue and Red strategy and game value to Blue for a notional case with $a = 5$, $c = 5$, $b = .5$ and $PkH = .1$


Red final strategy = .09

48 0403 RAND

We here plot the change in Red and Blue strategy over time as we approach the end of a conflict. The bottom line in the graph shows the optimum Red strategy, the percentage of SAMs that fire, as a function of the number of days remaining in the conflict. At the far left of the chart, 50 days before the end of the conflict, Red rarely fires SAMs. As the end of the conflict approaches, Red's firing rate increases, finally reaching 9 percent at the end of the war. The top line shows the optimum Blue strategy, the percentage of aircraft flying strike. Blue flies almost entirely strike missions 50 days from the end of the war, but it gradually increases the resources devoted to SEAD as the conflict draws to a close. The strategies seem to mirror each other, with Red firing more SAMs and Blue flying more SEAD as both sides approach the end of the war.

The middle line shows the value to Blue of each day's combat. Blue's value decreases as Red's SAM firing rate increases. However, Red cannot increase its firing rate prematurely, for fear that it would not have adequate reserves to last to the end of the war, leaving Blue in a very favorable situation. If Red is foolish enough to engage prematurely, Red will lose assets, which will make its situation much worse in the future.

Red Has Limited SAMs and High-Value Targets to Defend



In this case, $a = 1$, $c = 1$, and $b = 5$. That is, the targets are worth more than either the aircraft or the SAMs. This results in the high (71%) SAM engagement rate at the end of the war.

However, until the last 20 engagements, Red's optimal strategy is still to fire rarely, less than 10% of the time.

49 0403 RAND

We now repeat the calculation for a case in which Red should, in fact, fire fairly aggressively at the end of the conflict. The value of the struck targets is so high that Red should fire 71 percent of the time, but this result applies only to the last engagement of the conflict. When we work out the numbers, we find that until the last 20 engagements of the conflict Red should fire less than 10 percent of the time. The effect of the time horizon is strong. It forces Red to a conservative firing doctrine, even when the value of the targets is so high that Red would fire aggressively in a short war.

Note that this analysis is entirely from the point of view of operational commanders trying to achieve well-defined operational goals, i.e., strike targets, down aircraft, or kill SAMs. The analysis would be different if one included war termination conditions. For example, one might posit that downing a U.S. aircraft in the first week of the war would lead to war termination on terms favorable to Red, but that downing a U.S. aircraft two months into the war would not. In such a case, the value of an aircraft would not be constant throughout the conflict and the analysis would be different. The game theoretic and dynamic programming techniques described here are perfectly suited for such an analysis. A large body of work addresses game theory and war termination as well as war deterrence.⁵ We do not present such analysis here.

⁵ A classic text is T. C. Schelling *The Strategy of Conflict*, Harvard University Press, Cambridge, MA, 1960.

What If the War Duration Is Uncertain?

What if Red and Blue have different views as to the length of the war?
What if Blue expects a short war while Red expects a longer war?

Blue will invest more in SEAD than Red expects.

Red will fire SAMs less frequently than Blue expects.

This will not change the expected value of the engagement.

(This is a highly counter-intuitive result of the minimax theorem.
Mistakes are not costly if the enemy cannot anticipate them.)

If one side is correct and understands the enemy's error, they can reap undeserved benefits.

If the war is really long but Blue thinks it will be short, Red can exploit that by reducing the SAM firing rate.


If the war is really short but Red thinks it will be long, Blue can exploit that by reducing the SEAD effort and putting more assets into strike.

50 0403 RAND

We consider the case in which the duration of the war is uncertain. Clearly, the side with the correct view of the war's duration will reap benefits. We note that if Red thinks the war will be longer than Blue expects, Red will tend to fire SAMs less aggressively than Blue expects. Conversely, if Blue expects a long war and Red expects a short war, Red will fire more aggressively than Blue expects.

Two Views of War Duration

Consider the case in which the optimum Red strategy was to fire 71% of the time on the last engagement. Suppose Blue believes that the war will last only one engagement while Red believes it will last ten.



51 0403 RAND

These charts illustrate the operational implications of differences of opinion about the duration of the war. We assume here that Blue expects the war to be over after one round of engagement and that Red expects the war to be over after ten rounds of engagement. For the underlying parameters of the conflict, we have taken the numbers from the previous chart.

The left chart illustrates Blue's view of the conflict. Blue expects a high level of SAM engagement and is surprised that Red's strategy is so conservative, firing only 20 percent of the time instead of the expected 71 percent of the time.

The right chart illustrates Red's view of the conflict. Red expects Blue to mainly conduct strike and is surprised that Blue devotes so many assets to SEAD.

The immediate result of both sides' analysis of the opening engagements of the conflict is, "It looks like the other guy is making a big mistake." If the leaders are intelligent, they will immediately ask, "Does the enemy know something I don't?" There is no straightforward way to use game theory or any other mathematical tool to predict what will happen in such a situation. The leader who first understands what is going on will have the advantage, because he will be able to manipulate the opponent's perception. This is an area where the analysis of the game theorist must take a subordinate position to the intuition of the commander. Everyone on each side must desperately hope that their commander gets it right.

Resolution of Two Views of Duration

When two opponents have differing views of the duration of the conflict, they are designing strategies to different sets of MoEs. The condition that their interests are diametrically opposed no longer applies.

The short-term behavior of both sides is a complex problem in cooperative game theory. There is no one optimal solution. Bluffing can be effective. Even limited intelligence about the enemy's assessment of the situation can be very valuable.

Ultimately, the side that guessed wrong about the war's duration always achieves results less than or equal to the optimal solution. The side that guessed right always achieves optimal or better results.

52 0403 RAND

There are a few conclusions that one can proffer about situations in which the two sides hold conflicting views about the duration of the conflict. If the two sides really hold different views, game theory predicts that there is room for collaboration. It is possible to come up with results that look good to both sides. In technical language, this is cooperative game theory.

But it is phony cooperation. One side is right and one side is wrong. The side whose commander does not have a correct understanding of the duration of the conflict is at a terrible disadvantage. In situations like this, even limited and unreliable information can have great value.

Conclusions

Game theory provides a way to predict the behavior of an intelligent enemy confronted with U.S. military capabilities.

Successful U.S. capabilities, especially with respect to attacks on time critical targets, will often have the effect of causing the enemy to become paralyzed. The right move will be no move.

The right TCT systems for the United States are the ones that produce this effect. These are not necessarily the ones that produce the highest enemy casualties in traditional modeling, in which enemy CONOPs are fixed.

The anticipated duration of a conflict is very important to understanding the effectiveness of CONOPs.

53 0403 RAND

This document contains analysis of a variety of important military situations. The analysis is not of specific systems but rather is of a general nature, showing which considerations are important.

However, there is nothing about the techniques presented here that cannot be applied to real problems. The matrices might be larger and the arithmetic might be more complex, but the form of the analysis would be the same.

The main point is that game theory provides a mechanism for starting with the physics-based capabilities of weapon systems and calculating effects on enemy behavior. This mechanism is essential to calculating military effects from operational capabilities. Analysis that assumes particular enemy CONOPs will not provide insights into the effectiveness of U.S. capabilities that would logically induce an intelligent opponent to change its CONOPs. Of course, any highly effective U.S. system would usually induce exactly such a change of CONOPs in an intelligent enemy. Some mechanism for understanding the back-and-forth between U.S. capabilities and enemy actions is essential. Game theory provides such a mechanism.



OBJECTIVE ANALYSIS.
EFFECTIVE SOLUTIONS.

This product is part of the RAND Corporation documented briefing series. RAND documented briefings are based on research presented to a client, sponsor, or targeted audience in briefing format. Additional information is provided in the documented briefing in the form of the written narration accompanying the briefing charts. All RAND documented briefings undergo rigorous peer review to ensure that they meet high standards for research quality and objectivity. However, they are not expected to be comprehensive and may present preliminary findings. Major research findings are published in the monograph series; supporting or preliminary research is published in the technical report series.

DB-385-AF

RAND

Corporate Headquarters
1700 Main Street
P.O. Box 2138
Santa Monica, CA
90407-2138
TEL 310.393.0411
FAX 310.393.4818

Washington Office
1200 South Hayes Street
Arlington, VA 22202-5050
TEL 703.413.1100
FAX 703.413.8111

Pittsburgh Office
201 North Craig Street
Suite 202
Pittsburgh, PA 15213-1516
TEL 412.683.2300
FAX 412.683.2800

New York Office
215 Lexington Avenue
21st Floor
New York, NY 10016-6023
Council for Aid to Education
TEL 212.661.5800
New York External Affairs
TEL 212.661.3166
FAX 212.661.9766

RAND-Qatar Policy Institute
P.O. Box 23644
Doha, Qatar
TEL +974.492.7400
FAX +974.492.7410

RAND Europe—Berlin
Uhlandstrasse 14
10623 Berlin
Germany
TEL +49.30.310.191 0
FAX +49.30.310.19119

RAND Europe—Cambridge
Grafton House
64 Maids Causeway
Cambridge CB5 8DD
United Kingdom
TEL +44.1223.353.329
FAX +44.1223.358.845

RAND Europe—Leiden
Newtonweg 1
2333 CP Leiden
The Netherlands
TEL +31.71.524.5151
FAX +31.71.524.5191

www.rand.org

ISBN 0-8330-3259-3



9 780833 032591